DILEMMA DEL PRIGIONIERO RIPETUTO

Si consideri il dilemma del prigioniero, in cui 2 ladri vengono catturati, ma non vi sono sufficienti prove per condannarli. Vengono rinchiusi in celle separate e devono decidere se confessare (C) o non confessare (NC).
La struttura del gioco e' la seguente:

		B: NC	B: C
A	NC	1,1	5,0
	C	0,5	4,4

I numeri corrispondono ai mesi di prigione

Se il gioco viene giocato una sola volta, l'equilibrio di Nash e' (C,C).

Si supponga ora che lo stesso gioco venga giocato 2 volte. Il gioco viene risolto a ritroso.
Al tempo $t=2$ la soluzione del gioco e' (C,C).
Data questa soluzione, al tempo $t=1$ la matrice dei payoff diventa:

		B: NC	B: C
A	NC	5,5	9,4
	C	4,9	8,8

e l'equilibrio di Nash nel primo periodo e' (C, C).

Pertanto, in un dilemma del prigioniero ripetuto la soluzione del gioco e' confessare in ogni periodo, anche se non confessare garantirebbe una pena minore.

DECISIONI DI PRODUZIONE RIPETUTE UN NUMERO FINITO DI VOLTE.

Si supponga che due imprese debbano decidere quanto produrre a $t=1$ e $t=2$. In ogni periodo le decisioni sono simultanee. Ogni impresa può decidere se produrre una quantità alta (A) o bassa (B) del bene.
La matrice dei payoff è la seguente:

		Impresa 2: A	Impresa 2: B
Impresa 1	A	1, 1	5, 0
	B	0, 5	4, 4

Al tempo $t=2$ la soluzione di questo gioco è (A, A).
Dato questo equilibrio di Nash nel sottogioco $t=2$, la matrice dei payoff al tempo $t=1$ diventa:

		Impresa 2: A	Impresa 2: B
Impresa 1	A	2, 2	6, 1
	B	1, 6	5, 5

la cui soluzione è (A, A)

Pertanto, in entrambi i periodi le imprese producono una quantità alta del bene.

Si supponga ora che ciascuna impresa abbia la possibilità di produrre anche un livello medio del bene.
La matrice dei payoff diventa:

		Impresa 2		
		A	B	M
Impresa 1	A	1,1	5,0	0,0
	B	0,5	4,4	0,0
	M	0,0	0,0	3,3

Al tempo $t=2$ ci sono 2 equilibri di Nash: (A,A) ed (M,M).

Se (A,A) è l'equilibrio scelto a $t=2$, al tempo $t=1$ la matrice dei payoff diventa:

		Impresa 2		
		A	B	M
Impresa 1	A	2,2	6,1	1,1
	B	1,6	5,5	1,1
	M	1,1	1,1	4,4

e l'equilibrio nel primo periodo è anche (A,A) e (M,M).

Se (M,M) è l'equilibrio scelto nel secondo periodo, a $t=1$ la matrice dei payoff diventa:

		Impresa 2		
		A	B	M
Impresa 1	A	4,4	8,3	3,3
	B	3,8	7,7	3,3
	M	3,3	3,3	6,6

e anche in questo caso a $t=1$ l'equilibrio è (A,A) ed (M,M).

Come scegliere tra questi equilibri multipli?
Una strategia potrebbe essere: scegliere (M,M) nel primo periodo se anche (M,M) è scelto nel secondo periodo; altrimenti verrà scelto il livello di produzione alto nel secondo periodo.

Data questa strategia, la matrice dei payoff nel primo periodo diventa:

		Impresa 2		
		A	B	M
Impresa 1	A	2,2	6,1	1,1
	B	1,6	5,5	1,1
	M	1,1	1,1	6,6

che ha due equilibri: (A,A) ed (M,M).

Si supponga che a $t=1$ sia scelto l'equilibrio (A,A). Ma è credibile la strategia che anche a $t=2$ verrà scelto (A,A)?
Dato che a $t=2$ è disponibile la strategia (M,M), le imprese a $t=2$ potrebbero rinegoziare e scegliere (M,M), che garantisce un profitto maggiore, anche se nel primo periodo è stato scelto (A,A).
In tal caso, la matrice dei payoff a $t=1$ diventa:

		Impresa 2		
		A	B	M
Impresa 1	A	4,4	8,3	3,3
	B	3,8	7,7	3,3
	M	3,3	3,3	6,6

che ha 2 equilibri: (A,A) e (M,M)

Quindi, gli equilibri di questo gioco ripetuto sono: (A,A) a $t=1$ ed (M,M) a $t=2$; (M,M) a $t=1$ ed (M,M) a $t=2$.
La strategia di giocare (A,A) a $t=2$ se non si gioca (M,M) a $t=1$ non è quindi credibile
Si noti che se le imprese potessero cooperare nel primo periodo anche la strategia:

"giocare (M, M) nel secondo periodo se (B, B) viene scelto nel primo periodo; ma giocare (A, A) a $t=2$ se una scelta diversa da (B, B) venisse effettuata nel primo periodo" sarebbe disponibile.

Ma anche questa strategia non sarebbe credibile, per le stesse ragioni spiegate in precedenza.

Si noti che in quest'ultimo caso a $t=1$ sarebbe scelto (B, B), che non è un equilibrio di Nash del sottogioco.

GIOCO RIPETUTO UN NUMERO INFINITO DI VOLTE

Si supponga che il precedente gioco:

		impresa 2 A	impresa 2 B
Impresa 1	A	1, 1	5, 0
	B	0, 5	4, 4

sia ripetuto un numero infinito di volte.

Il valore attuale (V) di una sequenza infinita di pagamenti $\pi_1, \pi_2 \ldots$ è:

$$\pi_1 + \delta \pi_2 + \delta^2 \pi_3 + \ldots = \sum_{t=1}^{\infty} \delta^{t-1} \pi_t .$$

dove δ è il fattore di sconto che viene usato per riportare al presente i pagamenti futuri.

Si noti che δ può essere interpretato anche come il fattore di sconto quando il gioco può finire dopo un numero casuale di volte. In tal caso,

$$\delta = \frac{1-p}{1+r}$$ dove $1-p$ è la probabilità che il gioco continui il periodo successiv

Si consideri la seguente strategia:

"Cooperare nel primo periodo (B) ottenendo un profitto di 4, e continuare a cooperare nei periodi successivi a condizione che anche l'altra impresa cooperi. Se accade che in qualche periodo l'altra impresa sceglie A, giocare A all'infinito." (Trigger strategy).

Questa strategia è un equilibrio?

Se i giocatori cooperano nel primo periodo e continuano a cooperare nei periodi successivi, il valore attuale della sequenza dei pagamenti è:

$$4 + \delta 4 + \delta^2 4 + \cdots = \frac{4}{1-\delta}$$

Se un giocatore decidesse di produrre A nel primo periodo otterrebbe 5 in questo periodo ma 1 in tutti i periodi successivi. Il valore attuale di questi pagamenti sarebbe:

$$5 + \delta 1 + \delta^2 1 + \cdots = 5 + \delta(1 + \delta + \cdots)1 = 5 + \frac{\delta}{1-\delta}$$

[Si ricordi che $1 + \delta + \delta^2 + \cdots = \dfrac{1}{1-\delta}$]

Allora conviene cooperare se:

$$\frac{4}{1-\delta} \geq 5 + \frac{\delta}{1-\delta}$$

cioè se $\delta \geq 1/4$.

Quindi se il fattore di sconto è alto conviene giocare (B, B) in ogni periodo, anche se ciò non è un equilibrio di Nash del sottogioco.

Il fattore di sconto è alto se si attribuisce molta importanza ai pagamenti futuri e si pensa che la probabilità che il gioco continui sia alta.

Expected profit then becomes $2e^* - 2U_a - 2g(e^*)$, so the boss wishes to choose wages such that the induced effort, e^*, maximizes $e^* - g(e^*)$. The optimal induced effort therefore satisfies the first-order condition $g'(e^*) = 1$. Substituting this into (2.2.6) implies that the optimal prize, $w_H - w_L$, solves

$$(w_H - w_L) \int_{\varepsilon_j} f(\varepsilon_j)^2 \, d\varepsilon_j = 1,$$

and (2.2.8) then determines w_H and w_L themselves.

2.3 Repeated Games

In this section we analyze whether threats and promises about future behavior can influence current behavior in repeated relationships. Much of the intuition is given in the two-period case; a few ideas require an infinite horizon. We also define subgame-perfect Nash equilibrium for repeated games. This definition is simpler to express for the special case of repeated games than for the general dynamic games of complete information we consider in Section 2.4.B. We introduce it here so as to ease the exposition later.

2.3.A Theory: Two-Stage Repeated Games

Consider the ~~Prisoners' Dilemma~~ given in normal form in Figure 2.3.1. Suppose two players play this simultaneous-move game twice, observing the outcome of the first play before the second play begins, and suppose the payoff for the entire game is simply the sum of the payoffs from the two stages (i.e., there is no

Player 2

Player 1		L_2	R_2
	L_1	$1,1$	$5,0$
	R_1	$0,5$	$4,4$

Figure 2.3.1.

Player 2

Player 1		L_2	R_2
	L_1	$2,2$	$6,1$
	R_1	$1,6$	$5,5$

Figure 2.3.2.

discounting). We will call this repeated game the two-stage Prisoners' Dilemma. It belongs to the class of games analyzed in Section 2.2.A. Here players 3 and 4 are identical to players 1 and 2, the action spaces A_3 and A_4 are identical to A_1 and A_2, and the payoffs $u_i(a_1, a_2, a_3, a_4)$ are simply the sum of the payoff from the first-stage outcome (a_1, a_2) and the payoff from the second-stage outcome (a_3, a_4). Furthermore, the two-stage Prisoners' Dilemma satisfies the assumption we made in Section 2.2.A: for each feasible outcome of the first-stage game, (a_1, a_2), the second-stage game that remains between players 3 and 4 has a unique Nash equilibrium, denoted by $(a_3^*(a_1,a_2), a_4^*(a_1,a_2))$. In fact, the two-stage Prisoners' Dilemma satisfies this assumption in the following stark way. In Section 2.2.A we allowed for the possibility that the Nash equilibrium of the remaining second-stage game depends on the first-stage outcome—hence the notation $(a_3^*(a_1,a_2), a_4^*(a_1,a_2))$ rather than simply (a_3^*, a_4^*). (In the tariff game, for example, the firms' equilibrium quantity choices in the second stage depend on the governments' tariff choices in the first stage.) In the two-stage Prisoners' Dilemma, however, the unique equilibrium of the second-stage game is (L_1, L_2), regardless of the first-stage outcome.

Following the procedure described in Section 2.2.A for computing the subgame-perfect outcome of such a game, we analyze the first stage of the two-stage Prisoners' Dilemma by taking into account that the outcome of the game remaining in the second stage will be the Nash equilibrium of that remaining game—namely, (L_1, L_2) with payoff $(1,1)$. Thus, the players' first-stage interaction in the two-stage Prisoners' Dilemma amounts to the one-shot game in Figure 2.3.2, in which the payoff pair $(1,1)$ for the second stage has been added to each first-stage payoff pair. The game in Figure 2.3.2 also has a unique Nash equilibrium: (L_1, L_2). Thus, Dilemma is (L_1, L_2) in the first stage, followed by (L_1, L_2) in the second stage. Cooperation—that is, (R_1, R_2)—cannot be achieved in either stage of the subgame-perfect outcome.

This argument holds more generally. (Here we temporarily depart from the two-period case to allow for any finite number of repetitions, T.) Let $G = \{A_1, \dots, A_n; u_1, \dots, u_n\}$ denote a static game of complete information in which players 1 through n simultaneously choose actions a_1 through a_n from the action spaces A_1 through A_n, respectively, and payoffs are $u_1(a_1, \dots, a_n)$ through $u_n(a_1, \dots, a_n)$. The game G will be called the *stage game* of the repeated game.

Definition *Given a stage game G, let G(T) denote the* ***finitely repeated game*** *in which G is played T times, with the outcomes of all preceding plays observed before the next play begins. The payoffs for G(T) are simply the sum of the payoffs from the T stage games.*

Proposition *If the stage game G has a unique Nash equilibrium then, for any finite T, the repeated game G(T) has a unique subgame-perfect outcome: the Nash equilibrium of G is played in every stage.*[13]

We now return to the two-period case, but consider the possibility that the stage game G has multiple Nash equilibria, as in Figure 2.3.3. The strategies labeled L_i and M_i mimic the Prisoners' Dilemma from Figure 2.3.1, but the strategies labeled R_i have been added to the game so that there are now two pure-strategy Nash equilibria: (L_1, L_2), as in the Prisoners' Dilemma, and now also (R_1, R_2). It is of course artificial to add an equilibrium to the Prisoners' Dilemma in this way, but our interest in this game is expositional rather than economic. In the next section we will see that infinitely repeated games share this multiple-equilibria spirit even if the stage game being repeated infinitely has a unique Nash equilibrium, as does the Prisoners' Dilemma. Thus, in this section we

[13]Analogous results hold if the stage game G is a dynamic game of complete information. Suppose G is a dynamic game of complete and perfect information from the class defined in Section 2.1.A. If G has a unique backwards-induction outcome, then $G(T)$ has a unique subgame-perfect outcome: the backwards-induction outcome of G is played in every stage. Similarly, suppose G is a two-stage game from the class defined in Section 2.2.A. If G has a unique subgame-perfect outcome, then $G(T)$ has a unique subgame-perfect outcome: the subgame-perfect outcome of G is played in every stage.

	L_2	M_2	R_2
L_1	1,1	5,0	0,0
M_1	0,5	4,4	0,0
R_1	0,0	0,0	3,3

Figure 2.3.3.

analyze an artificial stage game in the simple two-period framework, and thereby prepare for our later analysis of an economically interesting stage game in the infinite-horizon framework.

Suppose the stage game in Figure 2.3.3 is played twice, with the first-stage outcome observed before the second stage begins. We will show that there is a subgame-perfect outcome of this repeated game in which the strategy pair (M_1, M_2) is played in the first stage.[14] As in Section 2.2.A, assume that in the first stage the players anticipate that the second-stage outcome will be a Nash equilibrium of the stage game. Since this stage game has more than one Nash equilibrium, it is now possible for the players to anticipate that different first-stage outcomes will be followed by different stage-game equilibria in the second stage. Suppose, for example, that the players anticipate that (R_1, R_2) will be the second-stage outcome if the first-stage outcome is (M_1, M_2), but that (L_1, L_2) will be the second-stage outcome if any of the eight other first-stage outcomes occurs. The players' first-stage interaction then amounts to the one-shot game in Figure 2.3.4, where $(3,3)$ has been added to the (M_1, M_2)-cell and $(1,1)$ has been added to the eight other cells.

There are three pure-strategy Nash equilibria in the game in Figure 2.3.4: (L_1, L_2), (M_1, M_2), and (R_1, R_2). As in Figure 2.3.2,

[14]Strictly speaking, we have defined the notion of a subgame-perfect outcome only for the class of games defined in Section 2.2.A. The two-stage Prisoner's Dilemma belongs to this class because for each feasible outcome of the first-stage game there is a unique Nash equilibrium of the remaining second-stage game. The two-stage repeated game based on the stage game in Figure 2.3.3 does not belong to this class, however, because the stage game has multiple Nash equilibria. We will not formally extend the definition of a subgame-perfect outcome so that it applies to all two-stage repeated games, both because the change in the definition is minuscule and because even more general definitions appear in Sections 2.3.B and 2.4.B.

	L_2	M_2	R_2
L_1	2,2	6,1	1,1
M_1	1,6	7,7	1,1
R_1	1,1	1,1	4,4

Figure 2.3.4.

Nash equilibria of this one-shot game correspond to subgame-perfect outcomes of the original repeated game. Let $((w,x),(y,z))$ denote an outcome of the repeated game—(w,x) in the first stage and (y,z) in the second. The Nash equilibrium (L_1, L_2) in Figure 2.3.4 corresponds to the subgame-perfect outcome $((L_1,L_2),(L_1,L_2))$ in the repeated game, because the anticipated second-stage outcome is (L_1, L_2) following anything but (M_1, M_2) in the first stage. Likewise, the Nash equilibrium (R_1, R_2) in Figure 2.3.4 corresponds to the subgame-perfect outcome $((R_1,R_2),(L_1,L_2))$ in the repeated game. These two subgame-perfect outcomes of the repeated game simply concatenate Nash equilibrium outcomes from the stage game, but the third Nash equilibrium in Figure 2.3.4 yields a qualitatively different result: (M_1, M_2) in Figure 2.3.4 corresponds to the subgame-perfect outcome $((M_1,M_2),(R_1,R_2))$ in the repeated game, because the anticipated second-stage outcome is (R_1, R_2) following (M_1, M_2). Thus, as claimed earlier, cooperation can be achieved in the first stage of a subgame-perfect outcome of the repeated game. This is an example of a more general point: if $G = \{A_1,\ldots,A_n;u_1,\ldots,u_n\}$ is a static game of complete information with multiple Nash equilibria then there may be subgame-perfect outcomes of the repeated game $G(T)$ in which, for any $t < T$, the outcome in stage t is not a Nash equilibrium of G. We return to this idea in the infinite-horizon analysis in the next section.

The main point to extract from this example is that credible threats or promises about future behavior can influence current behavior. A second point, however, is that subgame-perfection may not embody a strong enough definition of credibility. In deriving the subgame-perfect outcome $((M_1,M_2),(R_1,R_2))$, for example, we assumed that the players anticipate that (R_1, R_2) will be the second-stage outcome if the first-stage outcome is (M_1, M_2)

	L_2	M_2	R_2	P_2	Q_2
L_1	1,1	5,0	0,0	0,0	0,0
M_1	0,5	4,4	0,0	0,0	0,0
R_1	0,0	0,0	3,3	0,0	0,0
P_1	0,0	0,0	0,0	$4,\frac{1}{2}$	0,0
Q_1	0,0	0,0	0,0	0,0	$\frac{1}{2},4$

Figure 2.3.5.

and that (L_1, L_2) will be the second-stage outcome if any of the eight other first-stage outcomes occurs. But playing (L_1, L_2) in the second stage, with its payoff of $(1,1)$, may seem silly when (R_1, R_2), with its payoff of $(3,3)$, is also available as a Nash equilibrium of the remaining stage game. Loosely put, it would seem natural for the players to renegotiate.[15] If (M_1, M_2) does not occur as the first-stage outcome, so that (L_1, L_2) is supposed to be played in the second stage, then each player might reason that bygones are bygones and that the unanimously preferred stage-game equilibrium (R_1, R_2) should be played instead. But if (R_1, R_2) is to be the second-stage outcome after every first-stage outcome, then the incentive to play (M_1, M_2) in the first stage is destroyed: the first-stage interaction between the two players simply amounts to the one-shot game in which the payoff $(3,3)$ has been added to each cell of the stage game in Figure 2.3.3, so L_i is player i's best response to M_j.

To suggest a solution to this renegotiation problem, we consider the game in Figure 2.3.5, which is even more artificial than the game in Figure 2.3.3. Once again, our interest in this game is expositional rather than economic. The ideas we develop here to address renegotiation in this artificial game can also be applied to renegotiation in infinitely repeated games; see Farrell and Maskin (1989), for example.

[15]This is loose usage because "renegotiate" suggests that communication (or even bargaining) occurs between the first and second stages. If such actions are possible, then they should be included in the description and analysis of the game. Here we assume that no such actions are possible, so by "renegotiate" we have in mind an analysis based on introspection.

This stage game adds the strategies P_i and Q_i to the stage game in Figure 2.3.3. There are four pure-strategy Nash equilibria of the stage game: (L_1, L_2) and (R_1, R_2), and now also (P_1, P_2) and (Q_1, Q_2). As before, the players unanimously prefer (R_1, R_2) to (L_1, L_2). More importantly, there is no Nash equilibrium (x, y) in Figure 2.3.5 such that the players unanimously prefer (x, y) to (P_1, P_2), or (Q_1, Q_2), or (R_1, R_2). We say that (R_1, R_2) *Pareto-dominates* (L_1, L_2), and that (P_1, P_2), (Q_1, Q_2), and (R_1, R_2) are on the *Pareto frontier* of the payoffs to Nash equilibria of the stage game in Figure 2.3.5.

Suppose the stage game in Figure 2.3.5 is played twice, with the first-stage outcome observed before the second stage begins. Suppose further that the players anticipate that the second-stage outcome will be as follows: (R_1, R_2) if the first-stage outcome is (M_1, M_2); (P_1, P_2) if the first-stage outcome is (M_1, w), where w is anything but M_2; (Q_1, Q_2) if the first-stage outcome is (x, M_2), where x is anything but M_1; and (R_1, R_2) if the first-stage outcome is (y, z), where y is anything but M_1 and z is anything but M_2. Then $((M_1, M_2), (R_1, R_2))$ is a subgame-perfect outcome of the repeated game, because each player gets $4 + 3$ from playing M_i and then R_i but only $5 + 1/2$ from deviating to L_i in the first stage (and even less from other deviations). More importantly, the difficulty in the previous example does not arise here. In the two-stage repeated game based on Figure 2.3.3, the only way to punish a player for deviating in the first stage was to play a Pareto-dominated equilibrium in the second stage, thereby also punishing the punisher. Here, in contrast, there are three equilibria on the Pareto frontier—one to reward good behavior by both players in the first stage, and two others to be used not only to punish a player who deviates in the first stage but also to reward the punisher. Thus, if punishment is called for in the second stage, there is no other stage-game equilibrium the punisher would prefer, so the punisher cannot be persuaded to renegotiate the punishment.

2.3.B Theory: Infinitely Repeated Games

We now turn to infinitely repeated games. As in the finite-horizon case, the main theme is that credible threats or promises about future behavior can influence current behavior. In the finite-horizon case we saw that if there are multiple Nash equilibria of the stage game G then there may be subgame-perfect outcomes of the repeated game $G(T)$ in which, for any $t < T$, the outcome of stage t is not a Nash equilibrium of G. A stronger result is true in infinitely repeated games: even if the stage game has a unique Nash equilibrium, there may be subgame-perfect outcomes of the infinitely repeated game in which no stage's outcome is a Nash equilibrium of G.

We begin by studying the infinitely repeated Prisoners' Dilemma. We then consider the class of infinitely repeated games analogous to the class of finitely repeated games defined in the previous section: a static game of complete information, G, is repeated infinitely, with the outcomes of all previous stages observed before the current stage begins. For these classes of finitely and infinitely repeated games, we define a player's strategy, a subgame, and a subgame-perfect Nash equilibrium. (In Section 2.4.B we define these concepts for general dynamic games of complete information, not just for these classes of repeated games.) We then use these definitions to state and prove Friedman's (1971) Theorem (also called the Folk Theorem).[16]

Suppose the Prisoners' Dilemma in Figure 2.3.6 is to be repeated infinitely and that, for each t, the outcomes of the $t - 1$ preceding plays of the stage game are observed before the t^{th} stage begins. Simply summing the payoffs from this infinite sequence of stage games does not provide a useful measure of a player's payoff in the infinitely repeated game. Receiving a payoff of 4 in every period is better than receiving a payoff of 1 in every period, for example, but the sum of the payoffs is infinity in both cases. Recall (from Rubinstein's bargaining model in Section 2.1.D) that the discount factor $\delta = 1/(1 + r)$ is the value today of a dollar to be received one stage later, where r is the interest rate per stage. Given a discount factor and a player's payoffs from an infinite

[16]The original Folk Theorem concerned the payoffs of all the Nash equilibria of an infinitely repeated game. This result was called the Folk Theorem because it was widely known among game theorists in the 1950s, even though no one had published it. Friedman's (1971) Theorem concerns the payoffs of certain subgame-perfect Nash equilibria of an infinitely repeated game, and so strengthens the original Folk Theorem by using a stronger equilibrium concept—subgame-perfect Nash equilibrium rather than Nash equilibrium. The earlier name has stuck, however: Friedman's Theorem (and later results) are sometimes called Folk Theorems, even though they were not widely known among game theorists before they were published.